Adaptation of Milagro Implicit Monte Carlo to the Roadrunner Hybrid Architecture

Timothy M. Kelley, Paul J. Henning, CCS-2

Implicit Monte Carlo (IMC) is a method of simulating thermal X-ray transport for nonlinear, time-dependent problems such as supernova explosions and inertial confinement fusion. Milagro is a set of radiation-only applications used to drive the Jayenne Project's IMC libraries for method validation and testing.

The Roadrunner supercomputer architecture increases a standard Opteron cluster floating point throughput by an order of magnitude with Cell Broadband Engine (Cell BE) processors. To study the feasibility of running IMC on Roadrunner, we adapted a subset of the Milagro applications and the underlying transport libraries to the Roadrunner hybrid architecture. We measured a substantial speedup of 5x on a representative test problem using prototype hardware.

In order to learn as much as possible about how to use the hybrid Roadrunner design, our study took two approaches. One started with the existing code base and evolved incrementally, first adapting the core particle transport routines to the Cell BE processor, then restructuring the full application (running on the Opterons) to offload particle transport to the Cell BE. A more revolutionary approach freely changed data structures and algorithms, packing data into work blocks that can be processed on any available computer resource—Cell BE or Opteron. Both approaches used interfaces that abstract the accelerator (Cell BE) and decouple the generation of particles (or work blocks) from processing them. These design considerations should enable the changes made for Roadrunner to apply to other emerging architectures as well as conventional clusters.

In Milagro's 3D wedge representation of 2D RZ geometry, we devised a problem consisting of a radiation source propagating through a complex pattern of thick and thin materials (Fig. 1). The thick materials were ratio-zoned to better resolve transport boundary layers. Although the two approaches were pursued independently, both achieved similar speedups on double bend; Fig. 2 shows

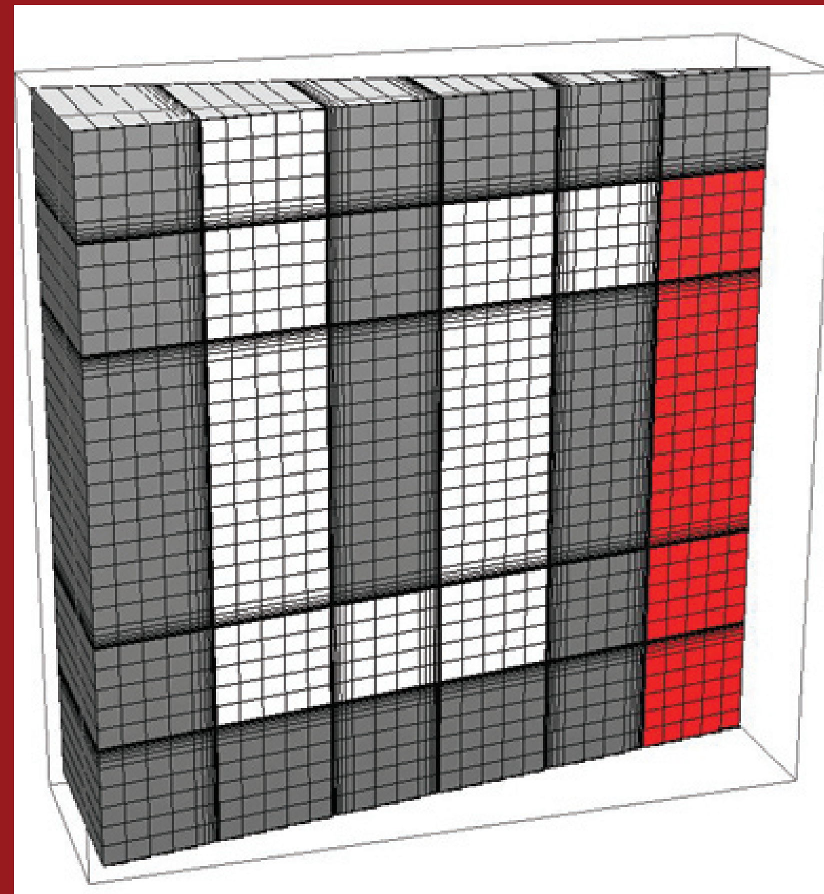

Fig. 1. The mesh used in the double bend test problem.

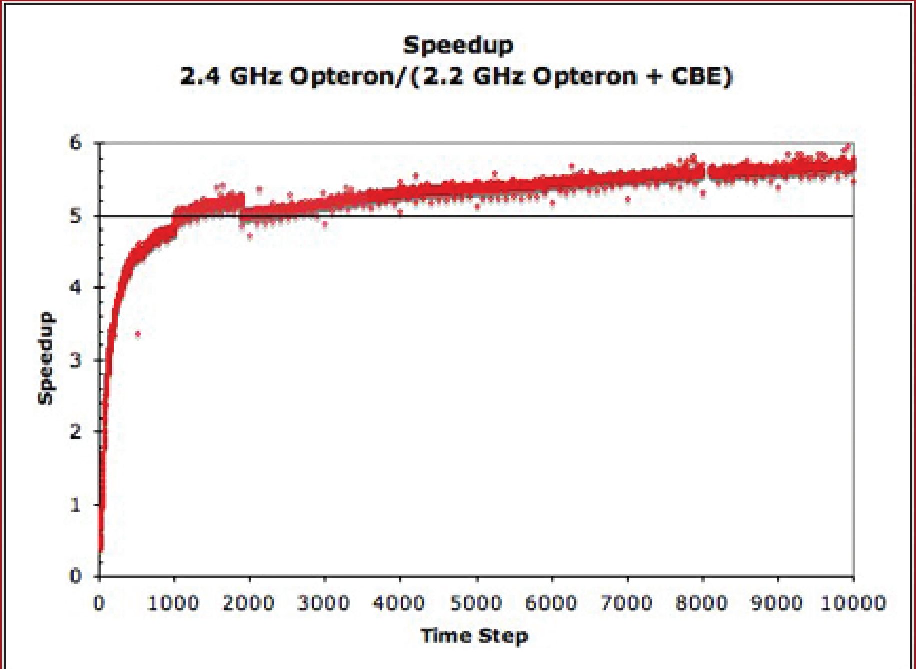

Fig. 2. Preliminary speedup of hybrid IMC on the double bend test problem as a function of time step, measured on prototype hardware. Very early in the simulation, there is not enough work to amortize the cost of running hybrid, but as the material heats up and more particles are used, hybrid quickly becomes more efficient than the single Opteron. By time step 2000, hybrid is five times faster (9.5 s wall clock hybrid, 48.0 s Opteron).

the speedup per time step. On prototype hardware—without Roadrunner's enhanced double precision Cell BE processors—we measured a 5x speedup over 10^4 time steps (1 μs simulation time). While the performance profile may vary for different problems, these results are excellent, and we are working on additional improvements.

In the future, we will merge the best of both approaches into the Jayenne project software base and research revolutionary transport methods enabled by this new architecture.

For more information contact Tim Kelley at tkelley@lanl.gov.

Funding Acknowledgments
- Department of Energy, National Nuclear Security Administration